arXiv:1608.04910v1 [stat.AP] 17 Aug 2016

Tweedie distributions for fitting semicontinuous health care utilization cost data

Christoph Kurz

August 18, 2016

Abstract

We explore a statistical distribution that can simultaneously model the probability of zero outcome for non-users of health care utilization and continuous costs for users. We compare this distribution to other commonly used models on example data and show that it fits cost data well and has some appealing properties that provide flexible use.

1 Introduction

In modeling cost data of health care utilization, the non-negative response variable is usually right-skewed with a positive probability of zero outcome for non-users. Such variables are called *semicontinuous* (10) and pose a number of problems: because of the point mass at zero, common models involving the Gamma or lognormal distributions have difficulty with such a mixture of discrete and continuous values. A popular way to account for this in the generalized linear models (GLM) framework are two-part models (3), which combine a binary model for the dichotomous event of having either zero or positive values with a continuous model for those having positive values. This complements a two-stage decision process which can be inadequate because the two decisions are usually not made independently (Winkelmann (15) and Van Ophem (14) discuss this for the case of physicians visits). Another more simple model, using a single distribution, is the Tobit model (13). This model is based on a zero-truncated Normal distribution but cannot handle *excess zeros*, i.e. the presence of more zeros in the data than what would be expected by the underlying distribution. In this linear regression setting, constant variance is assumed, which is also inadequate for cost data.

In this paper, we consider a single distribution GLM for cost data that can simultaneously model the zeros and continuous positives outcomes. The number of excess zeros can be arbitrarily high while still providing good support for the positive costs. Variance can be specified as some power of the mean. This model, based on the family of Tweedie densities (7) has already shown to

perform well in the case of rainfall precipitation (4) and insurance premium (11). To our knowledge, the Tweedie densities have not been used in health economic cost data modeling before. In the following we compare the Tweedie model to the Two-part (Binomial/Gamma) and Tobit regarding parameter estimates, model fit and prediction error. As analysts favor simple models that are easy to interpret, we restrict ourselves to these alternatives. For an overview of other, more specialized methods, we refer to (9).

The rest of this paper is structured as follows: Section 2 describes the properties of the Tweedie family of distributions and explains the proposed model; Section 3 outlines the data; Section 4 compares the Tweedie with the Two-part and Tobit model on these data; and Section 5 discusses the results.

Code and data to reproduce all analyses are available on the author's github page (1).

2 Tweedie Family Densities

We outline the model used in this paper as a special case of exponential dispersion models (EDMs) (7). This class of models is a broad family of distributions defined by the form

$$f(y|\theta,\phi) = a(y,\phi)\exp\left[\frac{y\theta-\kappa(\theta)}{\phi}\right],$$

where both the normalizing functions $a(\cdot)$ and $\kappa(\cdot)$ are known. θ is the natural parameter and $\phi > 0$ is called the dispersion parameter. Mean μ and variance of a random variable Y from an EDM are given by $E(Y) = \mu = \kappa'(\mu)$ and $\text{Var}(Y) = \kappa''(\mu)\phi$, respectively. The Tweedie family of distributions corresponds to special cases of EDMs where the power mean-variance relationship is characterized by $\text{Var}(\mu) = \phi\mu^p$ for $p \notin (0,1)$. The Tweedie family includes a number of familiar distributions, e.g. Normal ($p = 0$), Poisson ($p = 1$), Gamma ($p = 2$) and inverse Gaussian ($p = 3$).

For the case of cost data modeling the choice $p \in (1,2)$ is the most interesting one and our main focus here, because of its support for semicontinuous outcomes. Tweedie distributions in this range of p belong to the so-called compound Poisson-Gamma distributions (7). Let $M \sim \text{Pois}(\lambda)$ be a Poisson random variable, and let $X_i \overset{iid}{\sim} \text{Gamma}(\alpha,\beta)$ be Gamma distributed with $M \perp X_i$, then a random variable Z, defined by

$$Z = \begin{cases} 0, & \text{if } M = 0, \\ X_1 + X_2 + ... + X_M, & \text{if } M = 1,2,... \end{cases},$$

follows a compound Poisson-Gamma distribution, i.e. is a Poisson sum of Gamma random variables. If $M = 0$ then $Z = 0$, thus allowing for a probability mass at zero for non-users, where $\Pr(Z = 0) = \exp(-\lambda)$. If $M > 0$, then Z is the sum of M iid Gamma random variables, so conditional on M, $Z|M \sim \text{Gamma}(M\alpha, \beta)$, resulting in a continuous distribution for the positive outcome. With $M = m$ the distribution for $z > 0$ is therefore given as:

$$f(z|\lambda, \alpha, \beta) = \frac{\lambda^m \exp(-\lambda)}{m!} \frac{z^{m\alpha-1} \exp(-z/\beta)}{\beta^{m\alpha}\Gamma(m\alpha)}.$$

These parameters λ, α and β are related to the Tweedie distribution parameters μ, ϕ, and p by:

$$\lambda = \frac{\mu 2 - p}{\phi(2-p)}, \quad \alpha = \frac{2-p}{\alpha - 1}, \quad \beta = \phi(p-1)\mu^{p-1}.$$

Recovering the underlying marginal distribution of Z results in a non-closed form expression for the normalizing function $a(\cdot)$, based on Wright's generalized Bessel function $W(\cdot, \cdot, \cdot)$ (4, 16):

$$a(z, \phi) = \frac{1}{y} W(z, \phi, p).$$

Dunn and Smythe (6) show that this function is strictly convex and can be approximated by Stirling's formula for the Gamma function and a Fourier inversion method for the infinite series.

Because Tweedie distributions also belong to the exponential family of distributions, they can be used in the GLM framework (8). Besides the ability to model exact zeros and continuous outcomes, the idea that positive total costs are sums of smaller costs provide an intuitive appeal: Z is the total amount of expenses in a given period, M the number of utilization events, and X_i the expenses of the i-th event. In the following we show that the Tweedie distributions fit health care utilization cost data very well.

3 Data

We use data from the RAND Health Insurance Experiment (RAND HIE). This US study measured health care costs, among other outcomes, of people randomly assigned to different kinds of plans. Because of the random assignment, the reliability of health insurance coverage, and the availability of important variables for this application, these data provide an accurate base for cost modeling in our case.

As outcome variable we use the total costs, consisting of outpatient, drug,

supply, psychotherapy and inpatient expenses. We use covariates commonly considered to determine health care utilization. Among the socio-economic characteristics are age, gender, race, the logarithm of family income, LINC, the number of physical limitations, PHYSLM, the number of chronic diseases, DISEA, the logarithm of family size, LFAM, the education of household head in years, EDUCDEC, and a dummy variable indicating self-rated health as good, HLTHG. Insurance specific variables included are the log coinsurance rate plus 1, LOGC, a dummy for the individual deductible plan, IDP, the log of the participation incentive payment, LPI, and a maximum expenditure function, FMDE. A more detailed description of the data set and the variables can be found in (2).

We only use the first year of observation for each individual 18 years and older ($n = 3301$). There are 18.1% zero observations for the costs with a mean of 206.80 and a maximum of 17730.

4 Application

In this section, we apply the Two-part, the Tweedie and the Tobit model on the RAND HIE data. The Two-part model involves two estimations: first, we decide whether someone has zero or non-zero costs using a logistic regression. Second, conditional on having non-zero costs, we apply a Gamma model with log-link to the positive outcome. We also use the log-link for the Tweedie. The Tobit model features an underlying Normal distribution truncated at zero. For a more detailed description and justification of the Two-part and Tobit models see (10) and (9). The aim of this application is to show how model choice affects estimates, model fit and prediction in the case of semicontinuous health care cost data. No causal effects will be revealed in this study.

Table 1 presents the parameter estimation results for the discussed models. While the parameter estimates of the Tobit model are on a completely different scale, the Tweedie and the Gamma part of the Two-part model are very similar: all estimates share the same sign and have comparable values, leading to the same conclusions. Looking at the standard errors, we see that both Tweedie and Gamma estimations lead to very similar estimated standard errors, with the Gamma slightly underestimating the standard error. Furthermore, the log-likelihood of the Tweedie and Two-part are almost identical, suggesting comparable model fit: The Tweedie log-likelihood is -18874.4, while the Two-part has a combined log-likelihood of -18597.9 . The log-likehood of the Tobit model is slightly lower with a value of -21809.2. If we plot the true and estimated quantiles of the cost outcome for all distributions against each other, the Two-part has better model fit for the lower quantiles, while the Tweedie has slightly

higher support for upper quantiles. See Figure 1 for these Q-Q plots. This is likely because of the heavier tails of the Tweedie distribution. The Tobit model fits the central quantiles badly.

For the last comparison we randomly split the data into training set ($n = 2801$) and test set ($n = 500$). We build each model only on the training set and then evaluate on the test set. As comparison metric we chose root mean squared error (RMSE), defined by

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(\hat{y}_i - y_i)^2},$$

where $\hat{y}$ denotes the estimate and y the true value. Figure 2 plots the estimated total health care utilization costs against the true ones in the test set. Again, Two-part and Tweedie produce visually comparable estimates. The Tweedie model has the lowest RMSE with a value of 467.67, followed by the Two-part with 467.71 and the Tobit with 471.58.

The estimated value for the mean-variance power parameter in the Tweedie model is $p = 1.719$. Figure 3 shows the mean-variance plots for all 5% quantiles for the Tweedie model on the example data.

5 Discussion

This paper explores a single distribution GLM based on the Tweedie family of distributions for semicontinuous cost data. This model is comparable in model fit to the two-part Binomial/Gamma model but only includes a one-stage decision process, making it easier to interpret. The Tweedie model outperforms the Tobit model as the popular single distribution model for non-negative continuous data with a support for exact zeros. Thus, it provides an interesting alternative for modeling of health care utilization cost data as it has a natural support for cases where no utilization has occurred and for those it has. The theoretical justification is given since for the discussed case where the power parameter $p \in (1, 2)$, the Tweedie model can be explained as a Poisson sum of Gamma distributions. There, the number of utilization events is expressed by a Poisson distribution, and the amount of each utilization by a Gamma distribution.

While the theory of the Tweedie families is known for more than 20 years, only recently software packages that allow easy handling of these distributions became available (16, 5). Further research should explore the usefulness of Tweedie distributions with $p > 2$ as they provide similar shape to the Gamma but support heavier tails. Tweedie models in this range may be an attractive

alternative for the continuous part of a Two-part model or for cases without exact zeros and support a more flexible mean-variance relationship. The estimated mean-variance power parameter $p = 1.719$ may not appropriately reflect the true relationship. The fixed $p = 2$ in Gamma models is still too low, but values of p in the range of 2.2 to 2.3 seem more realistic.

Swallow et al (12) show in an ecological setting that a Bayesian hierarchical model based on the Tweedie densities provide further flexibility and removes this need to make strong assumptions about mean-variance relationships a priori. Such an hierarchical extension may further be useful to account for correlated effects by repeated measurement of individuals, for example in clinical trial settings or claims data.

Acknowledgements

The author thanks Rolf Holle and Irena Cenzer for helpful comments on the manuscript.

	Tobit		Tweedie		Two-part			
					Binomial		Gamma	
	Est.	SE	Est.	SE	Est.	SE	Est.	SE
Intercept	-212.276	118.391	4.253	0.543	-0.732	0.405	4.972	0.522
age	2.274	1.135	0.007	0.005	0.015	0.005	0.004	0.005
disea	6.319	1.695	0.015	0.008	0.044	0.008	0.008	0.007
physlm	214.712	34.484	0.688	0.150	0.321	0.156	0.664	0.145
logc	-23.994	17.530	-0.046	0.079	-0.184	0.065	-0.016	0.076
idp	-7.057	34.130	0.051	0.155	-0.061	0.130	0.057	0.149
lpi	-1.806	5.612	-0.028	0.025	0.005	0.022	-0.028	0.024
fmde	1.728	10.453	0.018	0.047	0.007	0.039	0.018	0.045
linc	18.113	11.951	0.066	0.055	0.096	0.038	0.032	0.055
lfam	-12.640	23.245	-0.013	0.106	0.054	0.088	-0.008	0.101
female	138.240	26.417	0.397	0.120	0.904	0.104	0.270	0.112
black	-166.908	39.181	-0.348	0.179	-1.040	0.128	-0.112	0.174
educdec	0.150	4.556	-0.011	0.021	0.044	0.018	-0.017	0.019
hlthg	-17.324	25.583	-0.067	0.116	0.102	0.100	-0.082	0.108
p			1.719					
dispersion	1		9.518		1		7.629	
log-likelihood	-21809.21		-18874.42		-1371.20		-17225.84	

Table 1: comparison of estimates

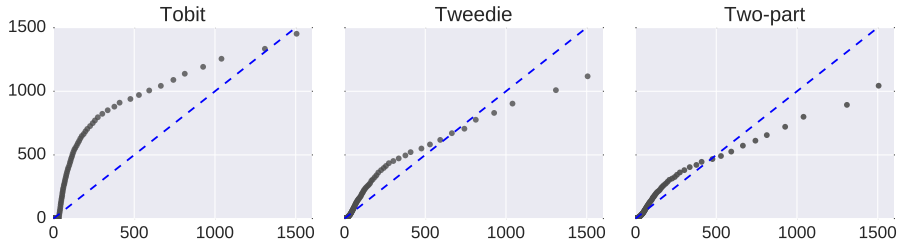

Figure 1: Q-Q plots for true and estimated quantiles of total health care utilization in the RAND HIE data for all models. Because of heavy outliers we do not show the last percentile. Quantile values closer to the dashed line represent a better match of empirical and estimated distributions.

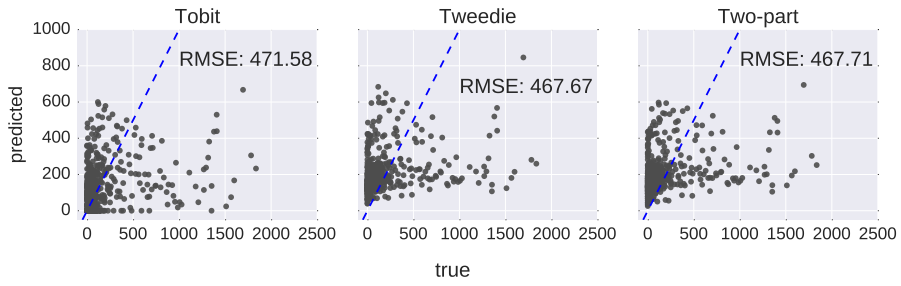

Figure 2: Plots of predicted versus true costs for total health care utilization for a test subset of the RAND HIE data. Predictions closer to the dashed line are more accurate. Some outliers exceeding the axis limits are not shown.

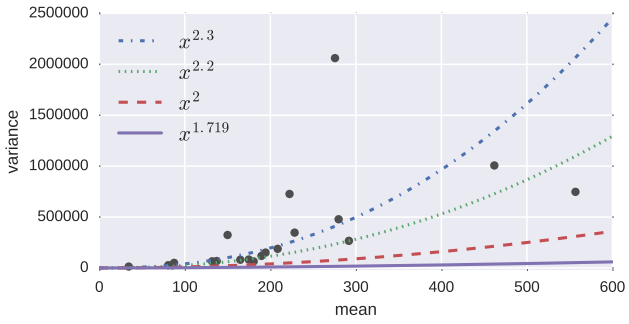

Figure 3: Mean-variance plots for all 5% quantiles for the Tweedie model. The solid line represents the estimated value for the mean-variance power parameter $p = 1.719$. Other values are plotted for comparison.

References

[1] `https://git.io/v6adW`. [Online; accessed 16-August-2015].

[2] P. Deb and P. K. Trivedi. The structure of demand for health care: latent class versus two-part models. *Journal of health economics*, 21(4):601–625, 2002.

[3] N. Duan, W. G. Manning, C. N. Morris, and J. P. Newhouse. A comparison of alternative models for the demand for medical care. *Journal of business & economic statistics*, 1(2):115–126, 1983.

[4] P. K. Dunn. Occurrence and quantity of precipitation can be modelled simultaneously. *International Journal of Climatology*, 24(10):1231–1239, 2004.

[5] P. K. Dunn. *tweedie: Tweedie exponential family models*, 2014. R package version 2.2.1.

[6] P. K. Dunn and G. K. Smyth. Evaluation of tweedie exponential dispersion model densities by fourier inversion. *Statistics and Computing*, 18(1):73–86, 2008.

[7] B. Jorgensen. *The theory of dispersion models*. CRC Press, 1997.

[8] P. McCullagh and J. A. Nelder. *Generalized linear models*, volume 37. CRC press, 1989.

[9] B. Mihaylova, A. Briggs, A. O'Hagan, and S. G. Thompson. Review of statistical methods for analysing healthcare resources and costs. *Health economics*, 20(8):897–916, 2011.

[10] Y. Min and A. Agresti. Modeling nonnegative data with clumping at zero: a survey. *Journal of Iranian Statistical Society*, 1(1):7–33, 2002.

[11] G. K. Smyth and B. Jørgensen. Fitting tweedie's compound poisson model to insurance claims data: dispersion modelling. *Astin Bulletin*, 32(01):143–157, 2002.

[12] B. Swallow, S. T. Buckland, R. King, and M. P. Toms. Bayesian hierarchical modelling of continuous non-negative longitudinal data with a spike at zero: An application to a study of birds visiting gardens in winter. *Biometrical Journal*, 2015.

[13] J. Tobin. Estimation of relationships for limited dependent variables. *Econometrica: journal of the Econometric Society*, pages 24–36, 1958.

[14] H. Van Ophem. The frequency of visiting a doctor: is the decision to go independent of the frequency? *Journal of Applied Econometrics*, 26(5):872–879, 2011.

[15] R. Winkelmann. Health care reform and the number of doctor visits—an econometric analysis. *Journal of Applied Econometrics*, 19(4):455–472, 2004.

[16] Y. Zhang. Likelihood-based and bayesian methods for tweedie compound poisson linear mixed models. *Statistics and Computing*, 23(6):743–757, 2013.